

Optimal Linear Precoding with Theoretical and Practical Data Rates in High-Speed Serial-Link Backplane Communication

Vladimir Stojanović^{1,2}, Amir Amirkhany¹ and Mark A. Horowitz¹

¹Department of Electrical Engineering, Stanford University, CA 94305, USA

²Rambus, Inc. Los Altos, CA 94022, USA

vlada@stanford.edu

Abstract—Multi-Gb/s high-speed links face significant challenges in keeping up with the increase in desired data rates. In the evaluation of achievable data rates, it is necessary to include both link-specific noise sources and implementation driven constraints. We construct models of these noise sources and constraints in order to estimate the theoretical limits of typical high-speed link channels. In order to estimate the data rates of practical baseband architectures, we solve the power constrained optimal linear precoding problem and formulate a bit-error rate (BER) driven optimization, including all link-specific noise sources. The problem is shown to be quasiconcave, hence, a globally optimal solution is guaranteed. Using this optimization framework, we show that practical data rates are mainly limited by inter-symbol interference (ISI) due to complexity constraints on the number of precoder and equalizer taps. After these constraints are removed, we further show that slicer resolution and sampling jitter are limiting the higher bandwidth utilization provided by multi-level modulations. Better circuits are needed to improve the bandwidth utilization to more than 2bits/dimension in baseband.

Keywords—precoding; capacity; serial-link; backplane; jitter

I. INTRODUCTION

Rapid increase in the bandwidth of optical links in Internet backbones (4x per generation) is driving the development of high-speed electrical links within the core of the router to sustain the increase in required switching throughput. Similar bandwidth growth occurs in other chip-to-chip applications. All these applications are entering the region of multi Gb/s rates over cables or backplane traces. Channel bandwidth limitations lead to Inter-Symbol Interference (ISI), which is becoming the dominant limit to the overall performance of the system. Both improved channels and chips that make better use of the available bandwidth are needed for link performance to continue to scale. However in many applications the channel is given, forcing a design to focus on how to achieve better bandwidth utilization.

Often tight power constraints and very high throughput requirements limit the complexity of digital communication algorithms that can be implemented in these chips, forcing designers to use less studied, inferior communication architectures. Furthermore, high-speed links have a set of noise

sources such as phase noise, sampling offset and supply noise that cannot be modeled as additive white Gaussian noise (AWGN) processes. This paper addresses these issues by describing a framework that can be used to evaluate the bandwidth limits of these high-speed channels, and provides information so that the designer can compare the efficiency of different communication architectures.

Electrical link channels are stationary and band-limited, but there is a large variation among different channels in the backplane due to different lengths and impedance mismatches (transmission lines, connectors, on-chip parasitics and terminations). Example backplane channels are shown in Fig. 1, illustrating the span in channel characteristics within one backplane (a), and between legacy and improved microwave engineered designs (b). The most dominant band-limiting effects are skin-effect and dielectric loss, while impedance mismatches cause multi-path-like reflections that form notches in the frequency response and persist in the channel many symbols after the main signal component has been received.

In order to accurately estimate the achievable data rates on these channels, it is essential to correctly model the link-specific noise sources, both fundamental and those that are a result of the high-throughput requirements. So we first focus on the derivation of these noise models. We then evaluate the impact of these noise sources on the capacity limits of the high-speed link channels, and the effective usable bandwidth.

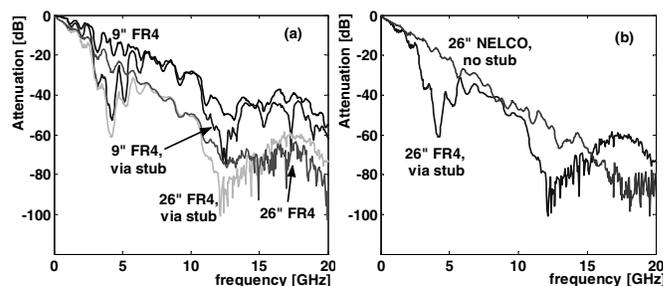


Figure 1. a) Frequency response of different channels within the same backplane: FR4 material, 9" and 26" trace length, top and bottom routing layers, b) Legacy and improved channels FR4 with stub and Nelco 26" counter bored to reduce reflections; (these two channels will be used for comparison throughout the paper)

Research supported by the MARCO Interconnect Focus Center and Rambus, Inc.

Practical signal-processing architectures in high-speed links are restricted to linear processing in the transmitter and range-restricted feedback equalization in the receiver, due to high-throughput requirements and lack of precision components (sampling resolution) in the receiver. Approximately optimal methods for linear precoder design under power constraints were first addressed in [1]. Building on their work, we formulate the precoding problem as a quasiconcave optimization problem with a globally optimal solution [2]. Given very low uncoded bit-error-rate (BER) requirements (lower than 10^{-15}), the optimization is BER driven with realistic noise sources for this type of environment, and peak power constrained to accommodate the integrated circuit (IC) signaling requirements.

Using this optimization framework we obtain the achievable data rates for practical baseband architectures with combinations of linear precoding and feedback equalization, allowing us to obtain data rates for system complexities similar to state-of-the-art links [3,4]. The optimization framework can easily be generalized to a multiple-input multiple-output (MIMO) system.

Given that energy efficiency is critical in high-speed links, we also estimate the energy costs of different architectures and how these costs scale with technology.

II. NOISE MODELS

Since noise, rather than deterministic effects like ISI, imposes capacity limits, we first derive the expressions for link-specific noise sources. While in many communication systems components can be designed so well that thermal noise is the real limiting factor, in high-speed link systems, high-throughput requirements yield circuits that result in non-negligible system noise, such as limited sampling resolution, sampling and carrier jitter.

A. Thermal Noise

The root causes of thermal noise in links are 50Ω terminations at the receiver. The noise figure of the receiver circuitry adds several dB to the termination noise level. Although the input bandwidth of the link is limited by the on-chip parasitics (due to transmit and receive circuits and electro-static discharge protection circuits), we assume that this bandwidth will always scale approximately with the signaling rate. In that scenario thermal noise spectral density is around $(1nV)^2/Hz$, which is $\sim 70dB$ down from the peak output energy of a typical IC transmitter at 10GHz Nyquist frequency, with transmitter output voltage swing constrained to $\pm 500mV$. This very high transmit signal-to-noise ratio (SNR) indicates that in order to truly estimate the performance of the link we need to consider other noise sources as well.

B. Slicer Resolution

Sampling resolution (the minimum voltage level that can be distinguished by the receiver slicer in the absence of other noise sources) is affected by several factors such as receiver static offset, input-referred supply noise and the overdrive required such that the slicer obtains the decision within a certain period of time (usually connected with system

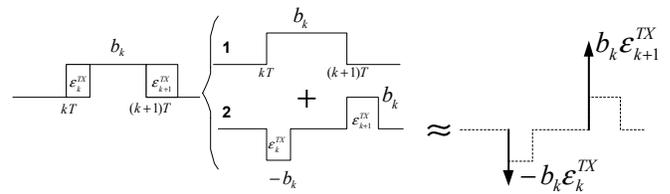


Figure 2. Jittered pulse decomposition. A symbol transmitted with jitter is converted to a symbol with no jitter (1), plus a noise term where the widths of the noise symbols (2) are equal to ϵ_k^{TX} and ϵ_{k+1}^{TX} .

throughput). Static offset occurs due to transistor mismatch from statistical process variations [5]. While it has statistical nature, the values are fixed once the chip is fabricated. While these offsets can be corrected to first order, this same mismatch limits the ability of the receiver to reject the supply noise of the IC environment [6]. The residual error is non-negligible, and we will use $\pm 10mV$ as the required sampling resolution. This value represents the residual error plus the required overdrive.

C. Sampling Jitter

The next noise source we consider is caused by jitter in the transmitter clock and the receiver clock. We need to map this jitter to voltage noise at the slicer input. Transmitter jitter modulates the position of both the beginning and the end of the symbol, while receiver jitter modulates the sampling position of the slicer.

Fig. 2 shows how we decompose a noisy symbol into a noiseless symbol l and two noise pulses¹ caused by the jitter 2 . Independence of the jitter process, ϵ , from the data stream, b , implies the independence of signals l and 2 , in Fig. 2. Since the two noise pulses are much narrower than the reference symbol pulse, we can approximate them with delta functions. When such noisy symbols pass through the channel filter, our approximation by delta functions is effectively equivalent to a zero-order approximation of the convolution integral.

With these assumptions, we create the system model, Fig. 3, where noiseless symbols pass through the standard channel pulse response block $p(nT)$, while the noise pulses pass through the impulse response block offset by half the symbol time $h(nT + T/2)$. The effect of receiver jitter is found using the

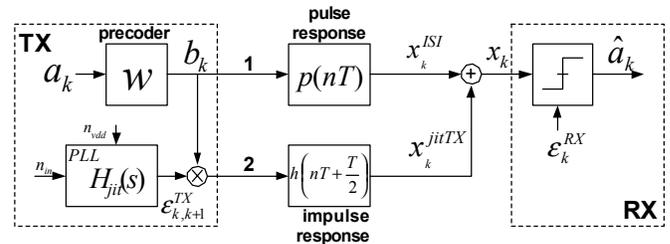


Figure 3. System model with transmitter and receiver sampling jitter. Since the noise pulses caused by transmitter jitter are narrow they are represented by impulses located at the edges of the symbol ($1/2$ a symbol from the symbol sample point) and after passing through the channel act as independent additive noise to the input data.

¹ In real high-speed link system implementations, basis functions are usually in the form of a square pulse as considered here, but the arguments are valid for any arbitrary pulse shape.

same model and shifting the entire transmit sequence by the amount of receive jitter.

The resulting expressions for samples at kT , (1a), corrupted by ISI, (1b), and voltage noise due to transmit and receive jitter, (1c,d) are:

$$x_k = x_k^{ISL} + x_k^{jitTX} + x_k^{jitRX} \quad (1a)$$

$$x_k^{ISL} = \sum_{n=-sbS}^{sbE} b_{k-n} p_n \quad (1b)$$

$$x_k^{jitTX} = \sum_{n=-sbS}^{sbE} b_{k-n} (h_{n-1} \epsilon_{k-n+1}^{TX} - h_n \epsilon_{k-n}^{TX}) \quad (1c)$$

$$x_k^{jitRX} = \epsilon_k^{RX} \sum_{n=-sbS}^{sbE} b_{k-n} (h_n - h_{n-1}) \quad (1d)$$

where ϵ_k^{TX} and ϵ_k^{RX} are samples of transmit and receive jitter, b_k the value of the transmitted symbol, with sbS and sbE as start and end indices of the impulse response sequence, and $p_n = p(nT)$ and $h_n = h(nT + T/2)$ samples of pulse and impulse responses, at nT and $nT + T/2$, respectively.

By deriving the autocorrelation functions for the noise samples from transmitter jitter, (1c) and receiver jitter, (1d), we were able to show that high frequency transmitter jitter creates much bigger voltage noise than high-frequency receiver jitter [7] (for example, with white Gaussian jitter with 1.4ps rms at 6.25GSymbol/s, the transmit and receive jitter induce noise voltages at the slicer input of 3mV and 1.6mV rms, respectively, out of +/- 500mV of transmitter output swing).

In case of transmit linear precoding, $b_k = \underline{w}^T \underline{a}_k$, where \underline{w} is the precoding vector and \underline{a}_k is the transmit alphabet vector, the autocorrelation functions of voltage noise from transmit and receive jitter can be shown to be of the form²:

$$\begin{aligned} R_{x^{jitTX,RX}}(m) &= \underline{w}^T E \left(\mathbf{A}(k) \underline{J}^{TX,RX}(k) \underline{J}^{TX,RX}(k+m)^T \mathbf{A}(k+m)^T \right) \underline{w} \\ \mathbf{A}(k) &= [\underline{a}_{k+sbS} \dots \underline{a}_{k-sbE}]; \quad \underline{a}_k = [a_{k+preW} \dots a_{k-postW}]^T \\ \underline{J}^{TX,RX}(k) &= [J_{-sbS}^{TX,RX}(k) \dots J_{sbE}^{TX,RX}(k)] \\ J_n^{TX}(k) &= \epsilon_{k-n+1}^{TX} h_{n-1} - \epsilon_{k-n}^{TX} h_n; \quad J_n^{RX}(k) = \epsilon_k^{RX} (h_n - h_{n-1}) \end{aligned} \quad (2)$$

where $\mathbf{A}(k)$ is the transmit alphabet matrix, and $preW$ and $postW$ are the number of taps before and after the main equalizer tap.

From (2), the autocorrelation can be compactly written as:

² Note that in this form, the variance of voltage noise due to transmit and receiver jitter ($R_{x^{jitTX,RX}}(0)$ sample of the autocorrelation in (2)) is actually a square of the l_2 norm in \underline{w} since the inner matrix is positive semi-definite. This convex form will be used later in Section V.

$$\begin{aligned} R_{x^{jitTX}}(m) &= \underline{w}^T \mathbf{S}_m^{TX} \underline{w}, \\ \mathbf{S}_m^{TX} &= E_a \sum_{j=-sbS}^{sbE} \sum_{k=-sbS}^{sbE} \mathbf{I}_{m+j-k} [h_{j-1} \ h_j] \begin{bmatrix} R_{\epsilon^{TX}}^{(m+j-k)} & -R_{\epsilon^{TX}}^{(m+j-k-1)} \\ -R_{\epsilon^{TX}}^{(m+j-k+1)} & R_{\epsilon^{TX}}^{(m+j-k)} \end{bmatrix} \begin{bmatrix} h_{k-1} \\ h_k \end{bmatrix} \end{aligned} \quad (3a)$$

$$\begin{aligned} R_{x^{jitRX}}(m) &= \underline{w}^T \mathbf{S}_m^{RX} \underline{w}, \\ \mathbf{S}_m^{RX} &= E_a R_{\epsilon^{RX}}(m) \sum_{j=-sbS}^{sbE} \sum_{k=-sbS}^{sbE} \mathbf{I}_{m+j-k} [h_{j-1} \ h_j] \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} h_{k-1} \\ h_k \end{bmatrix} \end{aligned} \quad (3b)$$

where E_a is the average energy of the transmit alphabet, a , $R_{\epsilon^{TX}}(m) = E(\epsilon_k^{TX} \epsilon_{k+m}^{TX})$ is the m^{th} sample of the autocorrelation function of transmit jitter, assuming that both jitter and data processes are stationary and that data is uncoded (i.e. independent), and \mathbf{I}_n is the identity matrix shifted right by n places.

D. Carrier Jitter (Phase Noise)

Carrier phase noise is present in any multi-tone implementation of high-speed links and induces some cross talk or interference between the real and imaginary parts of the signal. For narrow bandwidth communication, the received signal can be represented as a function of the transmitted symbol, channel frequency response and carrier phase noise in transmitter and receiver

$$x = H(j\omega_c t) a_{TX} e^{j\varphi_{TX}} e^{j(\theta_{noise}^{TX} - \theta_{noise}^{RX})} \quad (4)$$

where ω_c is the carrier frequency, $H(j\omega_c t)$ is the channel response at the carrier frequency, a_{TX} is the magnitude and φ_{aTX} the phase of the transmitted symbol, and θ_{noise}^{TX} and θ_{noise}^{RX} transmitter and receiver carrier phase noise, respectively.

The phase noise term in (4) results in mixing of the real and imaginary parts of the signal, which causes signal proportional noise with autocorrelation:

$$R_{x^{pn}}(m) = E_a \|H(j\omega_c t)\|^2 (R_{\theta_{noise}^{TX}}(m) + R_{\theta_{noise}^{RX}}(m)) \quad (5)$$

where E_a is the average transmit alphabet energy and $R_{\theta_{noise}^{TX}}(m)$ and $R_{\theta_{noise}^{RX}}(m)$ are the autocorrelation functions of transmitter and receiver carrier phase noise.

III. CAPACITY ESTIMATES

Having described the properties of the link-specific noise sources, we try to estimate their impact on the capacity of the link. To make the analysis more accurate, we also impose a peak power constraint, readily present in high-speed links.

In estimating the capacity of the link, we use the waterfilling solution³ [8]. Assuming the source signal distribution as

³ While this is exact for thermal noise, which is Gaussian, it is not exact for phase (carrier) noise since the capacity is achieved in that case when the sum of the signal and voltage noise due to phase noise is Gaussian. However, given that phase noise variance is usually much smaller than one, the Gaussian

Gaussian, for a fixed peak-to-average ratio (PAR), capacity achieved by waterfilling with $\Gamma=1$ (gap, defined in [9]) is a concave optimization problem⁴:

$$\lim_{N \rightarrow \infty} \left(\begin{aligned} &\text{maximize}_{E_n} \quad b = \frac{1}{2} \sum_{n=1}^N \log_2 \left(1 + \frac{E_n \|H_n\|^2}{\Gamma(\sigma_{thermal}^2 + E_n \|H_n\|^2 \sigma_{\theta n}^2)} \right) \\ &\text{s.t.} \quad \sum_{n=1}^N E_n = NE_{avg} = NE_{peak} PAR^{-1} \\ &\quad E_n \geq 0, \quad n=1, \dots, N \end{aligned} \right) \quad (6)$$

where $\sigma_{thermal}^2$ is thermal noise spectral density, $\sigma_{\theta n}^2$ is the sum of transmitter and receiver variances of phase noise of tone n , similar to (5), and N is the total number of tones.

The capacity curves with thermal noise are shown in Fig. 4a,b, for the best and worst channel, respectively. Due to the peak power constraint and very low BER requirements, we are interested in plotting capacity curves vs. clipping probability of the transmitted signal, determined by the PAR.

The phase noise of state-of-the-art frequency synthesizers, based on the LC oscillators [11], with a standard deviation of phase around 0.5° , further degrades the capacity by up to 5%. Using frequency synthesizers based on ring oscillators [3], and adding the phase noise from thermal and supply noise in carrier distribution buffers, results in up to 5° of phase noise standard deviation, which degrades the capacity by about 20%. In addition to this, with higher phase noise, capacity becomes less dependent on signal energy (and therefore clipping probability), since phase noise introduces a signal proportional noise source, (5).

It is also of practical interest to see how much un-coded integer constellation (like QAM) affects the data rate, so Fig. 5 shows integer loading curves with thermal noise for a gap of 13.3dB, corresponding to a BER of 10^{-15} . Modifying the Levin-Campello loading algorithm [12] to include the effects of carrier phase noise, we note that degradation in data rate is slightly more pronounced than that in capacity. With constellation increase, the required energy per channel increases faster than the minimum distance of the integer constellation points, causing the phase noise (which is

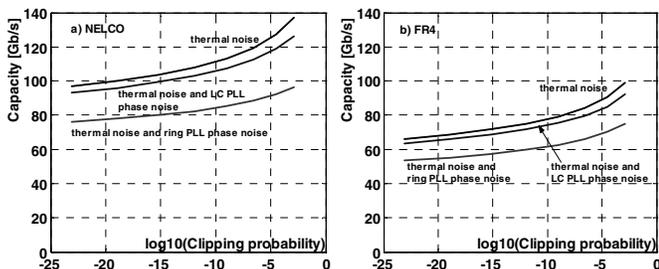


Figure 4. Capacity curves vs. clipping probability, for best (a) and worst channels (b), with thermal noise, phase noise from LC and ring oscillator based PLL.

distribution of the signal overwhelms the distribution of voltage noise due to phase noise and the resulting sum is mostly Gaussian.

⁴ This can be easily shown by examining the convexity in t on any energy line $E_n = E_{on} + tE_{sn}$ [10].

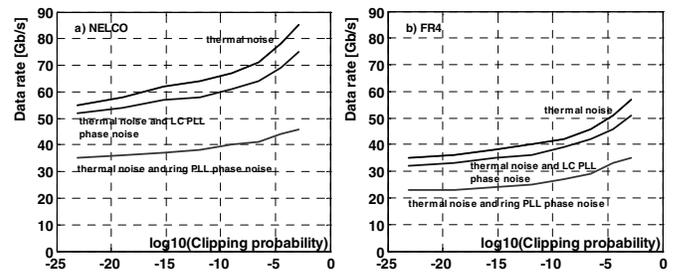


Figure 5. Data rate curves for integer loading with gap for BER= 10^{-15} , for best (a) and worst channels (b), with thermal noise, phase noise from LC and ring oscillator based PLL, obtained using the modified Levin-Campello loading algorithm [12], to account for the effects of phase noise.

proportional to energy), to become a more limiting factor than in the capacity case, where both signal and noise are proportional to energy and scale evenly.

The capacity estimates and data rate results in Figs. 4-5 show that even for very low BER requirements, and realistic noise sources, the achievable data rates are very high. Loading algorithms effectively use the channel up to Nyquist frequencies of 10GHz, as shown in Fig. 6, indicating that fundamental noise sources are too small to significantly limit the data rates.

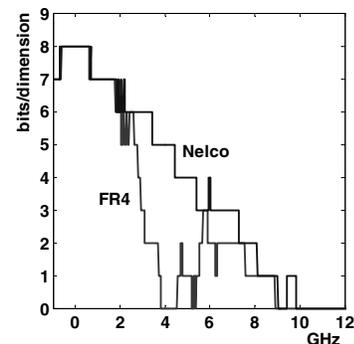


Figure 6. Integer loading with thermal noise.

Currently, state-of-the-art links achieve up to 10Gb/s rates over similar channels [3,4]. The main obstacle for practical systems achieving the data rates projected in Fig. 5 is effective ISI elimination. Additionally, great effort has to be put into improving the sampling resolution of the receiver, since this further reduces the achievable data rates.

In the following section we will illustrate the limitations of practical baseband techniques within the constraints imposed by the high-speed link IC environment.

IV. PRACTICAL ARCHITECTURES

A typical baseband high-speed link architecture is shown in Fig. 7, [3]. The link uses transmit precoding and analog decision feedback equalization. Transmit precoder replaces the standard receiver based feed-forward filter, since at very high symbol rates it is hard to create either a programmable analog feed-forward filter or high enough resolution (more than 6 bits) analog-to-digital converter [13] and implement the filter digitally.

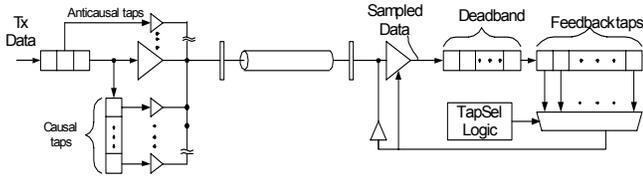


Figure 7. Baseband high-speed link architecture with transmit precoding and receiver decision feedback equalization

The decision feedback equalization (DFE) loop typically has a latency that is longer than a symbol time due to speed limitations in the receiver circuits, hence DFE is only used for reflection cancellation, [3]. The loop latency problem can be overcome to a limited extent using "loop unrolling" [14,15]. This enables the cancellation of a few immediate ISI taps that are essentially the most critical ones, since due to pulse dispersion they carry most of the ISI energy. Unfortunately, the complexity of loop unrolling grows as M^L with L being the number of taps and M the number of levels of modulation, hence its practical application is limited to PAM2 systems.

With increase in the desired data rates, the precoder and feedback equalizer lengths increase significantly, decreasing the power efficiency of the link. In addition, the precoding loss⁵ increases, limiting the achievable data rates in the presence of noise. In order to estimate the performance limits of such architectures, we derive a convex optimization framework that incorporates the link-specific noise sources in convex form and obtains globally optimal precoder and feedback filters.

V. OPTIMIZATION FRAMEWORK

It is well known that in the linear receiver equalizer problem, the minimization of mean square error (MSE), a quadratic form in equalizer taps, after unbiasing results in maximum signal-to-interference-and-noise ratio (SINR) and minimum BER⁶, [16].

Considering the system with a linear precoder, Fig. 8, we can formulate the MSE criterion, (7), whose form is similar to that of the linear receiver equalizer problem.

$$MSE(\underline{w}, g) = E_a \left(1 - 2g \underline{w}^T \mathbf{P} \mathbf{1}_\Delta + g^2 \underline{w}^T \mathbf{P} \mathbf{P}^T \underline{w} \right) + g^2 \sigma^2 \quad (7)$$

where \underline{w} is the precoding vector, \mathbf{P} is the Toeplitz matrix of the channel pulse response, g is the scalar receiver gain, $\mathbf{1}_\Delta$ is the

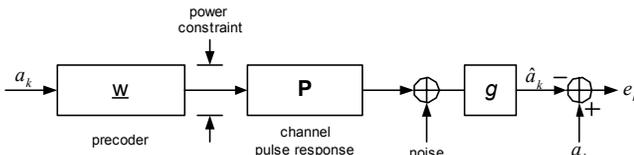


Figure 8. Precoding system with transmit power constraint and scalar gain in the receiver.

⁵ Similar to the linear receive equalizer noise amplification problem. While the l_2 norm of the linear equalizer amplifies noise in the receiver, the l_1 norm of the transmit precoder attenuates the transmitted signal in case of peak transmit power constraint.

⁶ BER here is defined assuming a mean distortion approximation, i.e. approximating the residual ISI as Gaussian noise.

system delay vector defined as $[0 \ 0 \ \dots \ 0 \ 1 \ 0 \ \dots \ 0]^T$ where one is in position $\Delta+1$ which represents system delay, E_a is the average energy of the transmitted alphabet a , and σ is the standard deviation of the AWGN source at the receiver.

Due to the power constraint, the precoder is not able to compensate the loss of the signal in the channel, but rather just compensate the ISI, while the gain element in the receiver compensates for the amplitude loss of the received signal. Thus the gain element effectively causes noise amplification.

In previous work on the optimization of a linear precoder, approximately optimal methods are derived, without using the gain element in the MSE criterion, [1]. It is also shown that such MSE criterion is sometimes inferior to the zero-forcing solution (ZFE), scaled to satisfy the power constraint. We extend that work by showing that the minimization of the MSE formulated using the receiver gain element, (7), is equivalent to SINR maximization, and therefore minimizes the BER.

From (7) the optimal gain g^* can be derived:

$$g^*(\underline{w}) = \frac{\underline{w}^T \mathbf{P} \mathbf{1}_\Delta}{\underline{w}^T \mathbf{P} \mathbf{P}^T \underline{w} + \sigma^2 / E_a} \quad (8)$$

which, when substituted in (7) yields

$$\frac{1}{SINR_{biased}} \triangleq \frac{MSE(\underline{w}, g^*)}{E_a} = 1 - \frac{E_a (\underline{w}^T \mathbf{P} \mathbf{1}_\Delta)^2}{E_a \underline{w}^T \mathbf{P} \mathbf{P}^T \underline{w} + \sigma^2} = \frac{1}{SINR_{unbiased} + 1} \quad (9)$$

where $SINR_{unbiased}$ represents the "true" (unbiased) signal-to-interference-and-noise ratio, and is defined as

$$SINR_{unbiased} \triangleq \frac{E_a (\underline{w}^T \mathbf{P} \mathbf{1}_\Delta)^2}{E_a \underline{w}^T \mathbf{P} (\mathbf{I} - \mathbf{1}_\Delta \mathbf{1}_\Delta^T) (\mathbf{I} - \mathbf{1}_\Delta \mathbf{1}_\Delta^T)^T \mathbf{P}^T \underline{w} + \sigma^2} \quad (10)$$

where $\underline{w}^T \mathbf{P} \mathbf{1}_\Delta$ represents the main tap of the received pulse response, and $\underline{w}^T \mathbf{P} (\mathbf{I} - \mathbf{1}_\Delta \mathbf{1}_\Delta^T) (\mathbf{I} - \mathbf{1}_\Delta \mathbf{1}_\Delta^T)^T \mathbf{P}^T \underline{w}$ the square of the l_2 norm of the residual ISI in the precoded pulse response.

The identity in (9) shows that minimization of the MSE defined as in (7), indeed results in maximization of the unbiased SINR. However, the convenient quadratic cost function is lost and the resulting problem is to maximize the SINR, which is a fractional quadratic programming problem, known to be non-convex [2,17].

Since our final goal is to minimize BER, starting from (10) directly, we note that the argument of the BER function is the square root of (10), resulting in the ratio of $\underline{w}^T \mathbf{P} \mathbf{1}_\Delta$, an affine function in \underline{w} , and the l_2 norm of $\underline{w}^T \mathbf{P} (\mathbf{I} - \mathbf{1}_\Delta \mathbf{1}_\Delta^T)$ and σ , which is convex in \underline{w} . It can be shown that maximizing this ratio is a quasiconcave programming problem with a global optimum [2], and can be efficiently solved by, for example, bisection [10].

Given that our final target is to minimize the actual BER, the BER function used in the optimization must be a very close approximation of the actual BER. Due to the very low BER requirements in high-speed links, it has been shown [18] that the Gaussian approximation of ISI, which leads to a BER

function defined as $Q(\sqrt{SINR_{unbiased}})$ is usually not very accurate for BERs $< 10^{-5}$. The main reason is that in fixed length precoders/equalizers ISI energy is dominated by a few very big residual components of dispersion component of ISI, but the total number of ISI taps is large due to reflections (which are much smaller than residual dispersion components). Since such ISI is not identically distributed, it cannot be well approximated with Gaussian distribution. To avoid this effect, we propose a mix of peak distortion and mean distortion criterions to achieve higher accuracy in BER approximation. It is only necessary to assume that a few big residual ISI taps are frequent enough to be considered as a constant shift from the mean value of the received signal, and the rest of the taps can be then well approximated with a Gaussian distribution. The resulting optimization, (11), is still quasiconcave.

$$\begin{aligned} \text{maximize } \gamma = & \frac{0.5d_{\min} \underline{w}^T \mathbf{P} \mathbf{I}_{\Delta} - V_{\text{peak}} \|\underline{w} \mathbf{P} \mathbf{I}_{PD}\|_1 - \text{offset}}{\left(E_a \underline{w}^T \mathbf{P} (\mathbf{I} - \mathbf{1}_{\Delta} \mathbf{1}_{\Delta}^T - \mathbf{I}_{PD}) (\mathbf{I} - \mathbf{1}_{\Delta} \mathbf{1}_{\Delta}^T - \mathbf{I}_{PD})^T \mathbf{P}^T \underline{w} + \sigma^2 \right)^{1/2}} \quad (11) \\ \text{s.t. } & \|\underline{w}\|_1 \leq 1 \end{aligned}$$

where the l_1 norm of \underline{w} is limited to 1 to satisfy the peak output power constraint, and \mathbf{I}_{PD} is the diagonal matrix that selects the residual ISI components to be considered for peak distortion. The average energy of the transmit alphabet, E_a , and minimum distance in transmit alphabet constellation, d_{\min} , assuming PAM modulation, are related to peak transmitter voltage V_{peak} by

$$d_{\min} = \frac{2V_{\text{peak}}}{M-1}, \quad E_a = T_{\text{symbol}} \frac{V_{\text{peak}}^2 (M+1)}{3(M-1)}, \quad M = 2^b. \quad (12)$$

Since the variances of voltage noise due to transmitter and receiver sampling jitter are convex (quadratic) functions of precoder taps, we can also add the impact of sampling jitter, from (3a,b), to the noise term σ^2 in (11), so the resulting noise variance is $\sigma^2 = \underline{w}^T \mathbf{S}_0^{TX} \underline{w} + \underline{w}^T \mathbf{S}_0^{RX} \underline{w} + \sigma_{\text{thermal}}^2$. The effect of limited slicer resolution is added to (11) in the term *offset*. In this way, we managed to include all of the described link-specific noise sources into the optimization framework in (11). The quasiconcave formulation of the optimization problem in (11) guarantees a globally optimal solution for the linear precoder, i.e. the one that achieves the minimum BER.

This framework can easily incorporate DFE in addition to transmit precoding. The optimal setting for feedback taps is to zero force the corresponding causal ISI in the received signal with precoding. Thus, prior to determining the precoder coefficients, we only need to pre-process the channel Toeplitz matrix, in such a way as to put "don't care" values on those residual ISI samples of the signal with precoding, whose indices correspond to the feedback taps that are to be used. This can be achieved by simply eliminating the columns of the channel Toeplitz matrix whose indices correspond to the time index of the feedback taps. Such a "punctured" Toeplitz matrix is then used in (11) to obtain the optimal transmit precoder coefficients. The feedback taps then just zero force the remaining response at particular tap indices.

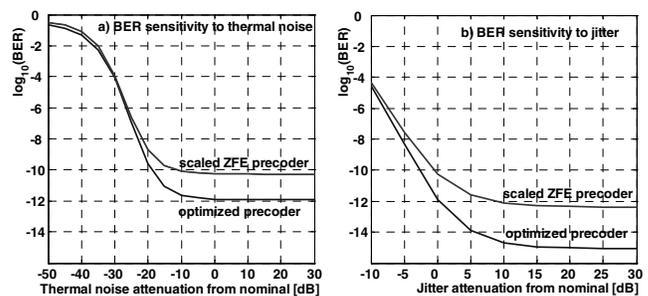


Figure 9. Sensitivity of BER to changes in: a) thermal noise, b) jitter, for 5 tap precoder with coefficients from scaled ZFE and optimization in (11). The system transmits PAM2 at 10Gb/s, with the Nelco channel.

Using this optimization framework we next evaluate the performance limits of practical implementations.

VI. DATA RATES FOR PRACTICAL IMPLEMENTATIONS

Fig. 9 gives the sensitivity of a 10Gb/s 2-PAM Nelco backplane link to thermal noise and jitter. It also compares the link results from using scaled ZFE and the optimization results from (11). The effect of changing the effective thermal noise is shown in Fig. 9a, and effective jitter in Fig. 9b. For Fig. 9, and all the data given in this section a noise figure of 7dB is added to $(1\text{nV})^2/\text{Hz}$ thermal noise of termination resistors, to account for thermal noise in the slicer. In addition when noted we will also assume 10mV of slicer resolution, and sampling jitter from a ring oscillator PLL with a standard deviation $\sigma_{\epsilon} = 5^\circ$ [3].

It is clear that noise from jitter is dominant in this link. Voltage noise due to jitter is especially harmful since it is proportional both to signal energy and jitter variance. This means that the only means to improve the system performance after system optimization is to minimize the jitter variance by careful circuit design.

Using our optimization framework we can now compare the expected performance of a number of different link architectures. Fig. 10 gives the achievable data rates if there were not any complexity constraints – using the precoder as a feedforward filter and assuming perfect feedback equalization in the receiver – and the links were limited by only thermal noise. The plot gives the performance range between the best and worst channels for different levels of modulation.

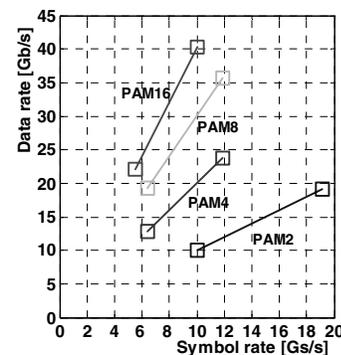


Figure 10. Data rates for 50-tap precoder with 80-tap feedback equalizer on best and worst channels with thermal noise, using different modulation levels. The lines correspond to the range of best achievable rates between the best and the worst channel. Target BER = 10^{-15} .

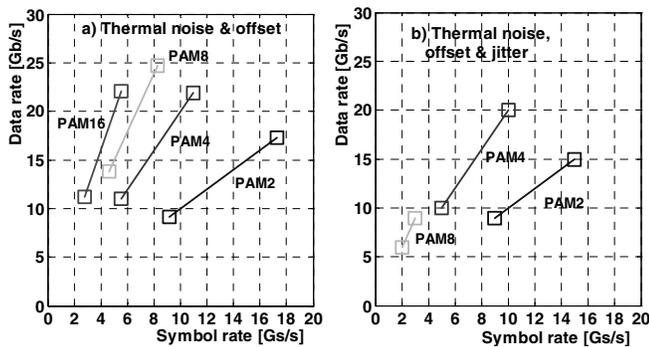


Figure 11. Data rates for 50-tap precoder with 80-tap feedback equalizer on best and worst channels, using different modulation levels in the presence of a) thermal noise and sampling resolution, b) thermal noise, sampling resolution and jitter.

Using higher levels of modulation in Fig. 10, the system more efficiently utilizes the effective channel bandwidth (9-12GHz, from bit loading in Fig. 6), and achieves very high data rates. However, we need to look at other sources of noise in order to evaluate the efficiency of multi-level modulations. In Fig. 11a, we add the receiver sampling resolution requirement, and in Fig. 11b, sampling jitter.

Sampling resolution imposes a constraint on the minimum distance between constellation points, so that one cannot add more constellation points within the peak power constraint without degrading system performance. This limits higher bandwidth utilization.

With good oscillator design, jitter noise is not dominant for small constellation sizes. However, since jitter noise energy is proportional to signal energy, it becomes more detrimental as energy remains the same and minimum distance between constellation points decreases. Therefore, jitter noise also prohibits the use of large constellations.

It is interesting to mention that a precoder filter alone has very poor performance, even without any constraint on complexity, due to the peak power constraint and large amount of ISI in the channels. Fig. 12 shows the projected data rates of practical baseband link architectures, keeping the complexity/power within the power budget of state-of-the-art links [3]. Since the large ISI can't be completely compensated now, higher PAM modulations start to fail.

From Figs. 10-12, we see that both the receiver resolution

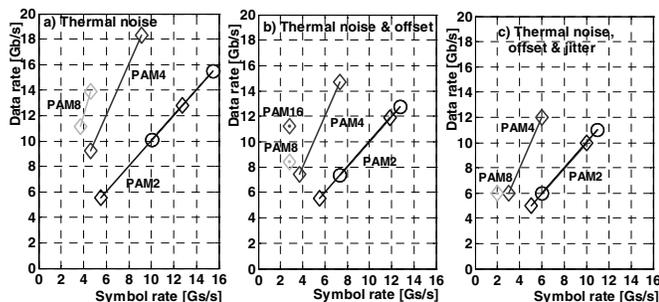


Figure 12. Achievable data rates with different noise sources for two architectures () 5 taps of transmit precoding with 20 taps of windowed reflection cancellation, similar to [3], with different levels of modulation, (o) same architecture, with "loop unrolling" by one extra tap of feedback equalization with no latency [7] (only PAM2 modulation is practical due to exponential growth in complexity).

and sampling jitter are limiting factors for the application of multi-level signaling techniques (higher than PAM4).

Any form of feedback equalization on dispersion ISI taps improves the performance, as shown in the PAM2 example where loop unrolling is used to cancel the first causal ISI tap. In order to achieve very low BERs, it is also essential to remove the long-latency reflections.

Our results clearly show that multi-level modulation together with precoding and feedback equalization with no latency is essential to achieving high data rates. In fact, the data rates of infinite length precoders and feedback equalizers are achievable with about 50 precoder taps and 80 feedback taps with no latency gaps. These rates, although high, are still not very close to the data rates projected in Fig. 5, for integer multi-tone constellations with thermal and phase noise. While improving the performance of baseband techniques is challenging, to achieve the rates projected in Fig. 5 will require the implementation of a practical multi-tone system that operates at these channel bandwidths.

VII. LINK ENERGY

Another important measure for high-speed links is the energy efficiency of data transmission, measured in mW/(Gb/s). Table I shows the energy efficiency cost for different link components that were taken from a recent link design [3]. A transmit precoding tap is more expensive than feedback equalization tap due to the larger size of the transmitter devices required to drive the desired output power. At the receiver, the size of the feedback taps can be smaller since the channel already attenuates the received signal. It is also interesting to note that although the cost of precoder and feedback equalizer taps increases with the number of levels of modulation (due to thermometer coding), the cost of the supporting blocks like synchronization (phase-locked loops and clock and data recovery loops) drops due to lower symbol rate requirements for the same data rate. It is also important to note that with ± 500 mV peak output swing in the transmitter, in a differential system the output power is fixed to 20mW, regardless of data rate.

TABLE I. ENERGY COST OF LINK COMPONENTS

PAM	TxTap	RxTap	RxSamp	PLL	CDR
2	1	0.3	2.2	8	11
4	1.5	0.45	5.9	4	5.5

All costs are in units of mW/(Gb/s). TxTap is cost per transmitter precoder tap, RxTap per feedback equalizer tap, RxSamp cost of sampling front-end, PLL cost of the phase-locked loop and CDR is the cost of the clock and data recovery loop [3].

Using the data from Table I, and achievable data rates for different architectures from Figs. 11 and 12, we plot in Fig. 13 the energy-efficiency of different architectures vs. data rate, for PAM2 and PAM4 modulations. The data indicates that for architectures with large numbers of taps, multi-level techniques are less energy efficient, since multi-level taps are more costly (due to thermometer coding), while for architectures with a small number of taps, multi-level architectures are more efficient since they decrease the amount of energy that is consumed in the supporting part of the link (for synchronization and clock generation).

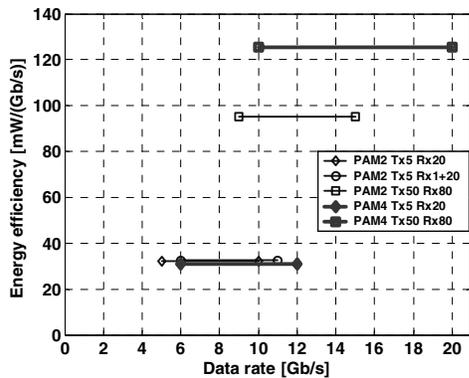


Figure 13. Energy efficiency of baseband architectures at different modulation levels.

The curves in Fig. 13 are given for current state-of-the-art $0.13\mu\text{m}$ CMOS technology. Conventional CMOS technology scaling assumes cubic energy scaling (quadratic in supply voltage and linear in capacitance). However, supply voltage scaling is severely limited in future technologies by transistor leakage, hence energy will most likely scale linearly with feature size due to capacitance scaling. This means that assuming much more complex filters will be possible in the near future might not be a good bet, and energy is likely to stay a key constraint in the future.

CONCLUSION

By providing models for link-specific noise sources, and showing how a power constrained precoder optimization can be formulated as a quasiconcave problem, we have been able to estimate both the capacity limit and the practical data rate limit for a number of backplane channels. While the effective bandwidth of these channels is limited to less than 10GHz, the capacity and integer constellation data rates are relatively high (40-100Gb/s). The problem is in achieving these high data rates in a practical system.

Our data indicates that both timing jitter and slicer sensitivity will seriously limit the data rate achievable in baseband to less than $\frac{1}{4}$ of capacity, even if one could do perfect DFE. Any practical equalization will be even worse, with the residual ISI becoming the largest source of noise in the system. In fact it seems like current links are running close to the maximum possible data rate for feasible baseband implementations.

Clearly, one way to get out of this box is to create a multi-tone system that can span the 10GHz available bandwidth. Our optimization framework can be extended to handle the MIMO analysis needed for this problem, and we are working on analyzing these systems next.

ACKNOWLEDGMENTS

The authors would like to acknowledge the help and support of R. Kollipara and B. Chia of Rambus for channel models, F. Chen, B. Garlepp and J. Zerbe of Rambus for link power data, D. Čabrić of UC Berkeley, and E. Alon of Stanford University for useful discussions. V. Stojanović and A. Amirkhany also thank I. Stojanović and B. Nezamfar for full-hearted help on the paper and technical discussions.

REFERENCES

- [1] B.R. Vojčić and W.M. Jang, "Transmitter Precoding in Synchronous Multiuser Communications," *IEEE Transactions on Communications*, vol. 46, no. 10, October 1998, pp. 1346-55.
- [2] S. Schaible, "Fractional programming," *Handbook of Global Optimization*, eds. Horst, R. and P.M. Pardalos, Nonconvex Optimization and its Applications, 2, Kluwer Academic Publishers, Dordrecht - Boston-London, 1995, pp. 495-608.
- [3] J. Zerbe *et al*, "Design, Equalization and Clock Recovery for a 2.5-10Gb/s 2-PAM/4-PAM Backplane Transceiver Cell," *IEEE International Solid-State Circuits Conference*, Feb. 2003, San Francisco.
- [4] R. Farjad-Rad *et al*, "0.622-8.0Gbps 150mW Serial IO Macrocell with Fully Flexible Preemphasis and Equalization", *IEEE Symposium on VLSI Circuits*, June 2003.
- [5] M.J. Pelgrom *et al*, "Matching properties of MOS transistors," *IEEE Journal Solid-State Circuits*, vol. 24, no. 5, pp. 1433-1439, Oct. 1989.
- [6] K-L.J. Wong and C-K.K. Yang, "Offset Compensation in Comparators with Minimum Input-Referred Supply Noise," *submitted to IEEE Journal Solid-State Circuits*.
- [7] V. Stojanović and M. Horowitz, "Modeling and Analysis of High-Speed Links," *IEEE Custom Integrated Circuits Conference*, September 2003.
- [8] R.G. Gallager, *Information Theory and Reliable Computation*, John Wiley & Sons, Inc., USA, 1968.
- [9] G.D. Forney Jr. and M. V. Eyuboglu, "Combined Equalization and Coding using Precoding," *IEEE Communications Magazine*, Dec. 1991, p.25-34
- [10] S. Boyd and L. Vandenberghe, *Convex Optimization*, book in preparation. (<http://www.stanford.edu/~boyd/cvxbook.html>)
- [11] H. Rategh, H. Samavati, and T. Lee, "A CMOS frequency synthesizer with an injection-locked frequency divider for a 5 GHz Wire LAN receiver," *IEEE J. Solid-State Circuits*, vol. 35, pp. 779-786, May 2000.
- [12] J. Campello, "Practical bit loading for DMT," *IEEE International Conference on Communications*, pp. 796-800, 1999.
- [13] K. Poulton *et al*, "A 20GS/s 8b ADC with a 1MB Memory in 0.18 μm CMOS," *IEEE International Solid-State Circuits Conference*, Feb. 2003, San Francisco.
- [14] K.K. Parhi, "High-Speed architectures for algorithms with quantizer loops," *IEEE International Symposium on Circuits and Systems*, vol. 3, May 1990, pp. 2357-2360
- [15] S. Kasturia and J.H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Selected Areas in Communications*, vol. 9, no. 5, Jun 1991, pp. 711-717.
- [16] J.M. Cioffi *et al*, "MMSE Decision-Feedback Equalizers and Coding-Part I: Equalization Results," *IEEE Transactions on Communications*, vol. 43, no. 10, October 1995, pp. 2582-2594.
- [17] V. Stojanović, G. Ginis and M.A. Horowitz, "Transmit Pre-emphasis for High-Speed Time-Division-Multiplexed Serial Link Transceiver," *IEEE International Conference on Communications*, pp. 1934 -1939, May 2002.
- [18] B. Ahmad, "Performance Specification of Interconnects," *DesignCon 2003*.